arXiv:1506.06983v1 [cs.DS] 23 Jun 2015

Linear Algorithms for Computing the *Lyndon Border Array* and the *Lyndon Suffix Array*

Ali Alatabbi [1], **Jacqueline W. Daykin** [1,2], **and M. Sohel Rahman** [3]

[1]Department of Informatics, King's College London, London, UK.
[2] Royal Holloway, University of London, UK,
[3]AℓEDA Group, Department of CSE, BUET, Dhaka-1000, Bangladesh

Abstract. We consider the problem of finding repetitive structures and inherent patterns in a given string $\boldsymbol{s}$ of length n over a finite totally ordered alphabet. A border $\boldsymbol{u}$ of a string $\boldsymbol{s}$ is both a prefix and a suffix of $\boldsymbol{s}$ such that $\boldsymbol{u} \neq \boldsymbol{s}$. The computation of the border array of a string $\boldsymbol{s}$, namely the borders of each prefix of $\boldsymbol{s}$, is strongly related to the string matching problem: given a string $\boldsymbol{w}$, find all of its occurrences in $\boldsymbol{s}$.
A *Lyndon word* is a primitive word (i.e., it is not a power of another word) which is minimal for the lexicographical order of its conjugacy class (i.e., the set of words obtained by cyclic rotations of the letters).
In this paper we combine these concepts to introduce the *Lyndon Border Array* $\mathcal{L}\beta$ of $\boldsymbol{s}$, whose i-th entry $\mathcal{L}\beta(\boldsymbol{s})[i]$ is the length of the longest border of $\boldsymbol{s}[1\,..\,i]$ which is also a Lyndon word. We propose linear-time and linear-space algorithms [1] for computing $\mathcal{L}\beta(\boldsymbol{s})$. Further, we introduce the *Lyndon Suffix Array*, and by modifying the efficient suffix array technique of Ko and Aluru [KA03] outline a linear time and space algorithm for its construction.

1 Introduction

Understanding complex patterns and repetitive structures in strings is essential for efficiently solving many problems in stringology [CR02]. For instance, Lyndon words are increasingly a fundamental and applicable form in the study of combinatorics on words [Lot83], [Lot05], [Smy03] - these patterned words have deep links with algebra and are rich in structural properties. Another important concept is a border $\boldsymbol{u}$ of a string $\boldsymbol{s}$ defined to be both a prefix and a suffix of $\boldsymbol{s}$ such that $\boldsymbol{u} \neq \boldsymbol{s}$. The computation of the border array of a string $\boldsymbol{s}$, that is of the borders of each prefix of $\boldsymbol{s}$, is strongly related to the string matching problem: given a string $\boldsymbol{w}$, find all of its occurrences in a string $\boldsymbol{s}$. It constitutes the "failure function" of the Morris-Pratt (1970) string matching algorithm [MP70].

Lyndon words were introduced under the name of *standard lexicographic sequences* [Lyn54,Lyn55] in order to construct a basis of a free Abelian group. Two

[1] The algorithms presented in this paper can be applied to computing the *co-Lyndon Border Array* by observing that a word $\boldsymbol{u} = u_1 u_2 \cdots u_n$ is a co-Lyndon word if the reversed word $\boldsymbol{u} = u_n u_{n-1} \cdots u_1$ is a Lyndon word.

strings are *conjugate* if they differ only by a cyclic permutation of their characters; a *Lyndon word* is defined as a (generally) finite word which is strictly minimal for the lexicographic order of its conjugacy class. For a non-letter Lyndon word $\boldsymbol{w}$, the pair $(\boldsymbol{u}, \boldsymbol{v})$ of Lyndon words such that $\boldsymbol{w} = \boldsymbol{uv}$ with $\boldsymbol{v}$ of maximal length is called the *standard factorization* of $\boldsymbol{w}$.

The set of Lyndon words permits the unique maximal factorization of any given string [CFL58,Lot83]. In 1983, Duval [Duv83] developed an algorithm for standard factorization that runs in linear time and space – the algorithm cleverly iterates over a string trying to find the longest Lyndon word; when it finds one, it adds it to the result list and proceeds to search in the remaining part of the string.

Lyndon words proved to be useful for constructing bases in free Lie algebras [Reu93], constructing de Bruijn sequences [FM78], computing the lexicographically smallest or largest substring in a string [AC95], succinct suffix-prefix matching of highly periodic strings [NS13]. Wider ranging applications include the Burrows-Wheeler transform and data compression [GS12], musicology [Che04], bioinformatics [DR04], and in relation to cryptanalysis [Per05]. Indeed the uses, and hence importance, of Lyndon words are increasing, and so we are motivated to investigate specialized Lyndon data structures.

The key contributions of this paper are as follows.

- By combining the important concepts of Lyndon words and borders of strings, we introduce here the *Lyndon Border Array* $\mathcal{L}\beta$ of $\boldsymbol{s}$, whose i-th entry $\mathcal{L}\beta(\boldsymbol{s})[i]$ is the length of the longest border of $\boldsymbol{s}[1\,..\,i]$ which is also a Lyndon word. We present an efficient linear time and space algorithm for computing the *Lyndon Border Array* $\mathcal{L}\beta$ for a given string (Section 4).
- In order to achieve the desired level of efficiency in the Lyndon Border Array construction we also present some interesting results related to Lyndon combinatorics, which we believe is of independent interest as well (Section 3).
- A complementary data structure, the *Lyndon Suffix Array*, which is an adaptation of the classic suffix array, is also defined; by modifying the linear-time construction of Ko and Aluru [KA03] we similarly achieve a linear construction for our Lyndon variant (Section 5). We also present a simpler algorithm to construct a Lyndon Suffix Array from a given Suffix Array (Section 5.1). The latter algorithm also runs in linear time and space.

2 Basic Definitions and Notation

Consider a finite totally ordered alphabet Σ which consists of a set of characters (equivalently letters or symbols). The cardinality of the alphabet is denoted by $|\Sigma|$.

A string (word) is a sequence of zero or more characters over an alphabet Σ. A string $\boldsymbol{s}$ of length $|\boldsymbol{s}| = n$ is represented by $\boldsymbol{s}[1\,..\,n]$, where $\boldsymbol{s}[i] \in \Sigma$ for $1 \leq i \leq n$. The set of all non-empty strings over the alphabet Σ is denoted by Σ^+, and the set of strings of length n by Σ^n. The empty string is the empty

sequence of characters (with zero length) denoted by ϵ, with $\Sigma^* = \Sigma^+ \cup \epsilon$; we write all strings in mathbold: $\boldsymbol{u}$, $\boldsymbol{v}$, and so on.

The i-th symbol of a string $\boldsymbol{s}$ is denoted by $\boldsymbol{s}[i]$, or simply s_i. We denote by $\boldsymbol{s}[i\,..\,j]$, or $s_i \cdots s_j$, the substring of $\boldsymbol{s}$ that starts at position i and ends at position j.

A string $\boldsymbol{w}$ is a substring, or factor, of $\boldsymbol{s}$ if $\boldsymbol{s} = \boldsymbol{uwv}$, where $\boldsymbol{u}, \boldsymbol{v} \in \Sigma^*$; specifically, a string $\boldsymbol{w} = w_1 \cdots w_m$ is a substring of $\boldsymbol{s} = s_1 \cdots s_n$ if $w_1 \cdots w_m = s_i \cdots s_{i+m-1}$ for some i. Words $\boldsymbol{w}[1\,..\,i]$ are called prefixes of $\boldsymbol{w}$, and words $\boldsymbol{w}[i\,..\,n]$ are called suffixes of $\boldsymbol{w}$. The prefix $\boldsymbol{u}$ (respectively suffix $\boldsymbol{v}$) is a proper prefix (respectively suffix) of a word $\boldsymbol{w} = \boldsymbol{uv}$ if $\boldsymbol{w} \neq \boldsymbol{u}, \boldsymbol{v}$.

For a substring $\boldsymbol{w}$ of $\boldsymbol{s}$, the string $\boldsymbol{uwv}$ for $\boldsymbol{u}, \boldsymbol{v} \in \Sigma^*$ is an extension of $\boldsymbol{w}$ in $\boldsymbol{s}$ if $\boldsymbol{uwv}$ is a substring of $\boldsymbol{s}$; $\boldsymbol{wv}$ for $\boldsymbol{v} \in \Sigma^*$ is the right extension of $\boldsymbol{w}$ in $\boldsymbol{s}$ if $\boldsymbol{wv}$ is a substring of $\boldsymbol{s}$; $\boldsymbol{uw}$ for $\boldsymbol{u} \in \Sigma^*$ is a left extension of $\boldsymbol{w}$ in $\boldsymbol{s}$ if $\boldsymbol{uw}$ is a substring of $\boldsymbol{s}$. Words that are both prefixes and suffixes of $\boldsymbol{w}$ are called borders of $\boldsymbol{w}$. By $border(\boldsymbol{w})$ we denote the length of the longest border of $\boldsymbol{w}$ that is shorter than $\boldsymbol{w}$.

A word $\boldsymbol{w}$ is periodic if it can be expressed as $\boldsymbol{w} = \boldsymbol{p}^k\boldsymbol{p}'$ where $\boldsymbol{p}'$ is a proper prefix of p, and $k \geq 2$. Moreover, a string is said to be primitive if it cannot be be written as $\boldsymbol{u}^k$ with $\boldsymbol{u} \in \Sigma^+$ and $k \geq 2$, i.e., it is not a power of another string. When $\boldsymbol{p}$ is primitive, we call it "the period" of $\boldsymbol{u}$. It is a known fact [CHL07] that, for any string $\boldsymbol{w}$, $per(\boldsymbol{w}) + border(\boldsymbol{w}) = |\boldsymbol{w}|$, where the period $p$er of a nonempty string is the smallest of its periods.

Definition 1. *(Border array) For a string $\boldsymbol{s} \in \Sigma^n$, the border array $\beta(\boldsymbol{s})[1\,..\,n]$ is defined by $\beta(\boldsymbol{s})[i] = |border(\boldsymbol{s}[1\,..\,i])|$ for $1 \leq i \leq n$.*

Proposition 2. *[MP70] The border of a string $\boldsymbol{s}$ (or the table $\beta(\boldsymbol{s})$ itself) can be computed in time $O(|\boldsymbol{s}|)$.*

A string $\boldsymbol{y} = \boldsymbol{y}[1\,..\,n]$ is a *conjugate* (or cyclic rotation) of $\boldsymbol{x} = \boldsymbol{x}[1\,..\,n]$ if $\boldsymbol{y}[1\,..\,n] = \boldsymbol{x}[i\,..\,n]\boldsymbol{x}[1\,..\,i-1]$ for some $1 \leq i \leq n$ (for $i = 1$, $\boldsymbol{y} = \boldsymbol{x}$). A *Lyndon word* is a primitive word which is minimal for the lexicographical order of its conjugacy class (i.e., the set of all words obtained by cyclic rotations of letters). Furthermore, a non-empty word is a Lyndon word if and only if it is strictly smaller in lexicographical order (lexorder) than any of its non-empty proper suffixes [Duv83,Lot83].

Throughout this paper, $\mathcal{L}$ will denote the set of Lyndon words over the totally ordered alphabet Σ, $\mathcal{L}_n$ will denote the set of Lyndon words of length n; hence $\mathcal{L} = \{\mathcal{L}_1 \cup \mathcal{L}_2 \cup \mathcal{L}_3, ..\}$. We next list several well-known properties of Lyndon words and border arrays which we later apply to develop the new algorithms.

Proposition 3. *[Duv83] A word $\boldsymbol{w} \in \Sigma^+$ is a Lyndon word if and only if either $\boldsymbol{w} \in \Sigma$ or $\boldsymbol{w} = \boldsymbol{uv}$ with $\boldsymbol{u}$, $\boldsymbol{v} \in \mathcal{L}$, $\boldsymbol{u} < \boldsymbol{v}$.*

Theorem 4. *[CFL58] Any word $\boldsymbol{w}$ can be written uniquely as a non-increasing product $\boldsymbol{w} = \boldsymbol{u}_1\boldsymbol{u}_2 \cdots \boldsymbol{u}_k$ of Lyndon words.*

Theorem 4 shows that there is a unique decomposition of any word into non-increasing Lyndon words ($\boldsymbol{u}_1 \geq \boldsymbol{u}_2 \geq \cdots \geq \boldsymbol{u}_k$).

Observation 1 *Let $\boldsymbol{\ell}$ be a Lyndon word ($\boldsymbol{\ell} \in \mathcal{L}$) where $\boldsymbol{\ell} = \ell_1\ell_2 \,..\, \ell_n$ (to avoid trivialities we assume $n > 1$), then*

(1) $border(\boldsymbol{\ell}) = 0$,
(2) $\ell_1 < \ell_n$,
(3) $\ell_1 \leq \ell_i | \ell_i \in \{\ell_2, .., \ell_{n-1}\}$,
(4) $\boldsymbol{\ell} < \ell_i \cdots \ell_n$, *for* $1 < i \leq n$.

Observation 2 *Given a string $\boldsymbol{s}$, then*

(1) $\beta(\boldsymbol{s})[1] = 0$,
(2) if $\boldsymbol{b}$ is a border of $\boldsymbol{s}$, and $\boldsymbol{b}'$ is a border of $\boldsymbol{b}$, then $\boldsymbol{b}'$ is a border of $\boldsymbol{s}$,
(3) $0 \leq \beta(\boldsymbol{s})[i+1] \leq \beta(\boldsymbol{s})[i] + 1$, *for* $1 \leq i < n$.

We now introduce the *Lyndon Border Array* and associated computation, illustrated in Example 7 below.

Definition 5. *(Lyndon Border Array) For a string $\boldsymbol{s} \in \Sigma^n$, the* Lyndon border array $\mathcal{L}\beta(\boldsymbol{s})[i]$ *is the length of the longest border of $\boldsymbol{s}[1\,..\,i]$ which is also a Lyndon word.*

Definition 6. *(Lyndon suffix array) For a string $\boldsymbol{s} \in \Sigma^n$, the* Lyndon Suffix Array *of $\boldsymbol{s}$ is the lexicographically sorted list of all those suffixes of $\boldsymbol{s}$ that form Lyndon words.*

Given a string $\boldsymbol{s}$ of length n, associated computational problems are: compute the Lyndon border and Lyndon suffix arrays; we address these problems in this paper.

Example 7. Consider the string $\boldsymbol{s} = abaabaaabbaabaab$. The following table illustrate the border array β of $\boldsymbol{s}$, the Lyndon border array $\mathcal{L}\beta$ of $\boldsymbol{s}$, the suffix array $\mathcal{A}$ of $\boldsymbol{s}$ and the Lyndon suffix array $\mathcal{L}S$ of $\boldsymbol{s}$.

i	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
$\boldsymbol{s}[i]$	a	b	a	a	b	a	a	a	b	b	a	a	b	a	a	b
$\beta[i]$	0	0	1	1	2	3	4	1	2	0	1	1	2	3	4	5
$\mathcal{L}\beta[i]$	0	0	1	1	2	1	1	1	2	0	1	1	2	1	1	2
$\mathcal{A}[i]$	5	13	2	10	6	14	3	11	0	7	15	4	12	1	9	8
$\mathcal{L}S[i]$	5	13	14	15	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1

3 Lyndon Combinatorics

This section introduces some new interesting combinatorial results on Lyndon words. In relation to the computation of the Lyndon Border Array, we here show how to find the shortest prefix of a string that is both border-free and not a Lyndon word. So assume that for a given string $\boldsymbol{s}$ of length n, we have $\boldsymbol{s}[1] = \gamma$. If $\boldsymbol{f_1}, \ldots, \boldsymbol{f_q}$ are factors of $\boldsymbol{s}$, we use $start(\boldsymbol{f_i})$ ($end(\boldsymbol{f_i})$) to denote the index of $\boldsymbol{f_i}[1]$ ($\boldsymbol{f_i}[|\boldsymbol{f_i}|]$) in $\boldsymbol{s}$, and say that j is an index of $\boldsymbol{f_i}$ if $start(\boldsymbol{f_i}) \leq j \leq end(\boldsymbol{f_i})$. An outline of the steps of the algorithm is as follows:

Algorithm *Shortest non-Lyndon Border-free Prefix (SNLBfP).*

1. Compute the Lyndon factorization of $\boldsymbol{s}$.
2. Apply binary search to find the first Lyndon factor $\boldsymbol{f_\mu}$ in the factorization starting with the largest letter μ which is strictly less than γ (if it exists).
3. Consider the maximal prefix $\boldsymbol{p}$ of $\boldsymbol{s}$ in which every factor $\boldsymbol{f_1}, \ldots, \boldsymbol{f_q}$ starts with γ; compute the border array $\beta(\boldsymbol{p})$ of $\boldsymbol{p}$.
4. Compute i, the smallest index of $\boldsymbol{p}$, such that $i > end(\boldsymbol{f_1})$ (i is not an index of $\boldsymbol{f_1}$) and $\beta(\boldsymbol{p})[i] = 0$;
   (a) if i does not exist then $i = end(\boldsymbol{f_q}) + 1$ (if it exists)
   (b) if $q = 1$ then $i = end(\boldsymbol{f_1}) + 1$ (if it exists).
5. Return $\boldsymbol{s}[1 \mathinner{..} i]$.

Claim 1 *Suppose $\boldsymbol{s}[1 \mathinner{..} i]$ is the shortest prefix of $\boldsymbol{s}$ that is both border-free and non-Lyndon. Then $i > end(\boldsymbol{f_1})$, i.e., Algorithm SNLBfP is correct in skipping the first Lyndon factor.*

Proof (Claim 1). The proof is by induction. Assume that the first Lyndon factor $\boldsymbol{f_1}$ is of length m (with $m \geq 2$ otherwise the Claim holds trivially). By Observation 1(1) and Observation 2(1) we have $\beta(\boldsymbol{f_1})[1] = \beta(\boldsymbol{f_1})[m] = 0$. The smallest j after 1 where $\beta(\boldsymbol{f_1})[j] = 0$ must index a Lyndon word $\boldsymbol{x}$ of the form $\boldsymbol{x} = \gamma^{j-1}\nu$, where $\nu > \gamma$. Hence, after the first letter (which is a Lyndon word), the next border-free prefix is also a Lyndon word.

Assume now that all border-free prefixes up to index t of $\boldsymbol{f_1}$ are Lyndon words (and hence nested), and suppose that the next border-free position is t'. Then we need to show that $\boldsymbol{y} = \boldsymbol{f_1}[1 \mathinner{..} t']$ is a Lyndon word; we proceed to show that $\boldsymbol{y}$ is less than each of its proper suffixes. Let $\boldsymbol{w_k} = \boldsymbol{f_1}[t' - k + 1 \mathinner{..} t']$ for $1 \leq k < t'$. Since $\boldsymbol{f_1}$ is a Lyndon word and minimal in its conjugacy class, then $\boldsymbol{f_1}[1 \mathinner{..} k] \leq \boldsymbol{w_k}$; further, since $\beta(\boldsymbol{f_1})[t'] = 0$ we cannot have equality and so $\boldsymbol{f_1}[1 \mathinner{..} k] < \boldsymbol{w_k}$ which implies that $\boldsymbol{y} < \boldsymbol{w_k}$ as required.

In other words, we have shown that any border-free prefix of $\boldsymbol{f_1}$ is a Lyndon word and the result follows.

□

Corollary 8. *Any border-free prefix of a Lyndon word is a Lyndon word.*

Lemma 9. Algorithm SNLBfP *is correct.*

Proof (Lemma 9). In Step 2 we identify the factor $\boldsymbol{f_\mu}$ which starts with the letter $\mu < \gamma$. From Lyndon principles, Observation 1(2),(3), it follows that no factor to the left of $\boldsymbol{f_\mu}$ contains the letter μ. Hence the prefix $\boldsymbol{s}[1\,..\,start(\boldsymbol{f_\mu})]$ is both border-free and non-Lyndon. However, it may not be the shortest one and so the algorithm continues to check through the prefix $\boldsymbol{p}$.

Now, consider the index i of $\boldsymbol{p}$ computed in Step 4. Suppose that i is an index of the factor $\boldsymbol{f_t}$; by Claim 1 we have $t > 1$. Let k be the length of the prefix $\boldsymbol{p_t}$ of $\boldsymbol{f_t}$ that ends at i, and $\boldsymbol{p_1}$ be the prefix of $\boldsymbol{f_1}$ of length k. By the Lyndon factorization we have $\boldsymbol{p_1} \geq \boldsymbol{p_t}$ (in lexicographic order). If $\boldsymbol{p_1} = \boldsymbol{p_t}$ then this contradicts $\beta(\boldsymbol{p})[i] = 0$. Hence $\boldsymbol{p_1} > \boldsymbol{p_t}$ and so the prefix $\boldsymbol{s}[1\,..\,i]$ is both border-free and not a Lyndon word. Hence Algorithm SNLBfP correctly returns $\boldsymbol{s}[1\,..\,i]$.

□

Lemma 10. Algorithm SNLBfP *runs in $O(n)$ time.*

Proof (Lemma 10). Step 1 can be computed in $O(n)$ time [Duv83,Lot05]. Step 2 applies an $O(\log n)$ binary search. In Step 3 we compute the border array of the prefix $\boldsymbol{p}$ of $\boldsymbol{s}$. Clearly Steps 3 and 4 can be completed in $O(n)$ time. Hence, the result follows.

□

Lemma 11. ***(Lyndon invalid point for border-free word)*** *Given a string $\boldsymbol{\ell} \in \Sigma^n$, $n > 1$, such that $border(\boldsymbol{\ell}) = 0$, if $\boldsymbol{\ell}$ is not a Lyndon word ($\boldsymbol{\ell} \notin \mathcal{L}$), then all the right extensions $\boldsymbol{\ell\ell'}$ of $\boldsymbol{\ell}$, such that $\boldsymbol{\ell'} \in \Sigma^+$, are not Lyndon words either.*

— We refer to this condition as the $\mathcal{L}$-fail condition and to the point (index) of where it occurs as the Lyndon invalid point.

Proof (Lemma 11). Since $\boldsymbol{\ell}$ is border-free ($border(\boldsymbol{\ell}) = 0$) but not a Lyndon word, then let $\boldsymbol{r}$ be the rotation of $\boldsymbol{\ell}$ which is the Lyndon word of the conjugacy class. Write $\boldsymbol{\ell} = \boldsymbol{pq}$ such that $\boldsymbol{r} = \boldsymbol{qp}$ and $\boldsymbol{qp} < \boldsymbol{pq}$.

Case (1) - If $|\boldsymbol{p}| = |\boldsymbol{q}|$, then since $\boldsymbol{r}$ is border-free and a Lyndon word, $\boldsymbol{q} \neq \boldsymbol{p}$. Further, since $\boldsymbol{qp} < \boldsymbol{pq}$, we have $\boldsymbol{q} < \boldsymbol{p}$. It follows that $\boldsymbol{q\ell'p} < \boldsymbol{pq\ell'}$ and so $\boldsymbol{\ell\ell'}$ cannot be a Lyndon word.

Case (2) - If $|\boldsymbol{q}| > |\boldsymbol{p}|$, then $\boldsymbol{r}[1..|\boldsymbol{p}|] < \boldsymbol{p}$ (i.e., $\boldsymbol{q}[1\,..\,|\boldsymbol{p}|] < \boldsymbol{p}$), we have $\boldsymbol{q} < \boldsymbol{p}$. It follows that $\boldsymbol{q\ell'p} < \boldsymbol{pq\ell'}$ and so $\boldsymbol{\ell\ell'}$ cannot be a Lyndon word.

Case (3) - If $|\boldsymbol{q}| < |\boldsymbol{p}|$. We have $\boldsymbol{r} = \boldsymbol{qp} \in \mathcal{L}$, where $\boldsymbol{q} = \boldsymbol{q_1 q_2} \cdots \boldsymbol{q_j}$, $\boldsymbol{p} = \boldsymbol{p_1 p_2} \cdots \boldsymbol{p_k}$. We are required to show that $\boldsymbol{r'} = \boldsymbol{pq\ell'} \notin \mathcal{L}$; so suppose that $\boldsymbol{r'} \in \mathcal{L}$. From $\boldsymbol{r}$ we have that $\boldsymbol{q}_1 \leq \boldsymbol{p}_1$, while from $\boldsymbol{r'}$ we have $\boldsymbol{p}_1 \leq \boldsymbol{q}_1$, which together implies $\boldsymbol{q}_1 = \boldsymbol{p}_1$. From the rotation of $\boldsymbol{r}$ starting $\boldsymbol{p}_1\boldsymbol{p}_2$ we have $\boldsymbol{q}_2 \leq \boldsymbol{p}_2$, while from the rotation of $\boldsymbol{r'}$ starting $\boldsymbol{q}_1\boldsymbol{q}_2$ we find that $\boldsymbol{p}_2 \leq \boldsymbol{q}_2$ giving $\boldsymbol{q}_2 = \boldsymbol{p}_2$. We continue this argument for $|\boldsymbol{q}|$ elements which shows that the given word $\boldsymbol{pq}$ has a border.

In the first two cases the order is decided within the first $|\boldsymbol{p}|$ elements.

□

Lemma 12. ***(Lyndon invalid point for bordered word)*** *Given a string $\boldsymbol{\ell} \in \Sigma^n$, $n > 1$, such that $border(\boldsymbol{\ell}) > 0$, let $\{\boldsymbol{\ell'^1}, \boldsymbol{\ell'^2}, \boldsymbol{\ell'^3} \cdots\}$ be right extensions of $\boldsymbol{\ell}$ by $\{1, 2, 3, \cdots\}$ characters in Σ respectively. Then $\boldsymbol{\ell}$ and all its right extensions are not Lyndon words if $\boldsymbol{\ell'^1}[|\boldsymbol{\ell'^1}|] < \boldsymbol{\ell'^1}[border(\boldsymbol{\ell}) + 1]$.*

— Similarly to Lemma 11, we refer to this condition as the $\mathcal{L}$-fail condition and to the point (index) of where it occurs as the Lyndon invalid point.

Proof (Lemma 12). The lemma follows from the following two cases:

1. $\boldsymbol{\ell} \notin \mathcal{L}$, this is immediate from the hypothesis that $border(\boldsymbol{\ell}) > 0$ and Observation 1(1).
2. Consider the right extension $\boldsymbol{\ell'^m}$ of $\boldsymbol{\ell}$, where $m \geq 1$. Then the suffix $\boldsymbol{\ell'^m}[n - border(\boldsymbol{\ell}) + 1 \,..\, n + m]$ of $\boldsymbol{\ell'^m}$ is lexicographically less than $\boldsymbol{\ell}'^m[1 \,..\, border(\boldsymbol{\ell}) + 1]$ and consequently less than $\boldsymbol{\ell'^m}$, contradicting the property that a Lyndon word is strictly smaller than any of its proper suffixes. Hence no right extension of $\boldsymbol{\ell}$ is a Lyndon word.

In case (2) the order is decided within the first $border(\boldsymbol{\ell}) + 1$ elements.

□

Fact 1 *For a given string $\boldsymbol{s} \in \Sigma^n$, suppose we have computed $\beta(\boldsymbol{s})$. Then, for $1 \leq i \leq n$, the following holds true:*

$$\mathcal{L}\beta(\boldsymbol{s})[i] \in \{\beta(\boldsymbol{s})[i],\ \beta(\boldsymbol{s})[\beta(\boldsymbol{s})[i]],\ \beta(\boldsymbol{s})[\beta(\boldsymbol{s})[\beta(\boldsymbol{s})[i]]],\ \ldots,\ 0\}.$$

A binary Lyndon word can also be expressed in terms of Lyndon properties of the integer parameters (exponents) given by its *Run Length Encoding*. For a binary string $\boldsymbol{\ell}$, let $\texttt{RLE}(\boldsymbol{\ell_a})$ denote the encoding (as a string) of the subsequence of $\boldsymbol{\ell}$ consisting of all letters a but no letter b $(p', p_1; \ldots; p_m)$, similarly $\texttt{RLE}(\boldsymbol{\ell_b})$ denotes the encoding of the subsequence of $\boldsymbol{\ell}$ consisting of all letters b $(q', q_1; \ldots; q_m)$.

Proposition 13. *For a given string $\boldsymbol{s} \in \Sigma^+$, the* Run Length Encoding *$RLE(\boldsymbol{s})$ and subsequently, the list $\mathcal{R}(\boldsymbol{s})$ can be computed in time and space linear in the size of the given string $\boldsymbol{s}$.*

Now we have the following Proposition to check if a given binary word is a Lyndon word or not (using the *Run Length Encoding* of the binary word).

Proposition 14. *Let $\mathcal{L}$ be the set of Lyndon words over an alphabet Σ, where $\Sigma = \{a, b\}$ and $a < b$. For a non-letter word $\boldsymbol{\ell} = \boldsymbol{\ell}_1\boldsymbol{\ell}_2 \cdots \boldsymbol{\ell}_n$ and its corresponding exponents list $\mathcal{R}(\mathcal{L})$, we have the following:*

Case $m = 0$**,** *we have $p_j, q_j = 0$ for $j \in \{1 \,..\, m\}$ then $\boldsymbol{\ell} \in \mathcal{L}$ if and only if $p', q' > 0$.*

Case $m > 0$**,** *we have $p', q', p_j, q_j > 0$ then*

1. *$\boldsymbol{\ell} \in \mathcal{L}$ if and only if $p' > p_j$ for $j \in \{1 \,..\, m\}$.*
2. *if $p' = p_m$ and $p' > p_j$ for $j \in \{1 \,..\, m - 1\}$ then $\boldsymbol{\ell} \in \mathcal{L}$ if and only if $q' < q_m \wedge \lambda < 0$.*

3. *if there exists* p_j *such that* $p' = p_j$ *and* $p' > p_m$ *then* $\boldsymbol{\ell} \in \mathcal{L}$ *if and only if* $q' \leq q_j \wedge \lambda < 0$ *for* $j \in \{1 \,..\, m-1\}$.

where λ *is the index of Lyndon invalid point i.e.* $\mathcal{L}$*-fail condition (As defined in Lemma. 11 and Lemma. 12).*

Lemma 15. *Let* $\boldsymbol{\ell}$ *be a binary word with* $\boldsymbol{\ell}[1] = a$, $\boldsymbol{\ell}[n] = b$ *and associated encodings* $\mathtt{RLE}(\boldsymbol{\ell_a})$ *and* $\mathtt{RLE}(\boldsymbol{\ell_b})$. *Then* $\boldsymbol{\ell}$ *is a Lyndon word if and only if, either*
(i) $\mathtt{RLE}(\boldsymbol{\ell_a})$ *is a Lyndon word on the alphabet* $\{1 > 2 > 3 > \cdots\}$, *or*
(ii) $\mathtt{RLE}(\boldsymbol{\ell_a})$ *is a repetition of a Lyndon word as in (i) and* $\mathtt{RLE}(\boldsymbol{\ell_b})$ *is a Lyndon word on the alphabet* $\{1 < 2 < 3 < \cdots\}$.

Proof (Lemma 20). First we consider necessity. So suppose that (i) holds. Consider any rotation $\boldsymbol{\ell_a^r}$ of $\boldsymbol{\ell_a}$ (including those with split runs of a's). Then $\mathtt{RLE}(\boldsymbol{\ell_a}) < \mathtt{RLE}(\boldsymbol{\ell_a^r})$ in lexorder over $\{1 > 2 > 3 > \cdots\}$. Now suppose that (ii) holds. Then for a rotation $\boldsymbol{\ell_a^r}$ of $\boldsymbol{\ell_a}$, either $\mathtt{RLE}(\boldsymbol{\ell_a}) < \mathtt{RLE}(\boldsymbol{\ell_a^r})$ in lexorder over $\{1 > 2 > 3 > \cdots\}$, or $\mathtt{RLE}(\boldsymbol{\ell_a}) = \mathtt{RLE}(\boldsymbol{\ell_a^r})$ and $\mathtt{RLE}(\boldsymbol{\ell_b}) < \mathtt{RLE}(\boldsymbol{\ell_b^r})$ in lexorder over $\{1 < 2 < 3 < \cdots\}$.

For sufficiency, the conditions guarantee that $\boldsymbol{\ell} \in \mathcal{L}$.

□

Lemma 20 shows that a binary string can be decomposed into its Lyndon factors with a double application of Duval's linear factorization algorithm [Duv83], first on the exponents of a and then on those for b – hence linear in $|\mathcal{R}(\boldsymbol{\ell})|$ (once $\mathcal{R}(\boldsymbol{\ell})$ is constructed in time linear in $|\ell|$). It follows that a word can be tested to be a Lyndon word in linear time: check whether there is more than one factor in the factorization.

Lemma 16. *Suppose we are given a binary string* $\boldsymbol{s}$ *of length* n *and the corresponding list of exponents* $\mathcal{R}(\boldsymbol{s})$. *Then we can compute the array* $\Psi(\boldsymbol{s})$ *in* $O(n)$ *time.*

Proof (Lemma 19). We first focus on computing the i-th element $\Psi[i](\boldsymbol{s})$ of the array $\Psi(\boldsymbol{s})$. Clearly, by Proposition 18, it takes constant number of comparisons to determine whether the prefix $\boldsymbol{s}[1 \,..\, i]$ is a Lyndon word or not. Since we have to repeat the steps in Proposition 18 for n prefixes, therefore $\Psi(\boldsymbol{s})$ can be done in linear time once $\mathcal{R}(\boldsymbol{s})$ is computed.

□

4 Lyndon Border Array Computation

In this section we develop an efficient algorithm for computing the Lyndon Border Array. We first recall an interesting relation that exists for borders which we refer to as the *Chain of Borders* henceforth. Since every border of any border of $\boldsymbol{s}$ is also a border of $\boldsymbol{s}$ (Observation 2(2)), it turns out that, the border array $\beta(\boldsymbol{s})$ compactly describes all the borders of every prefix of $\boldsymbol{s}$. For every prefix $\boldsymbol{s}[1 \,..\, i]$ of $\boldsymbol{s}$, the following sequence

$$\beta(\boldsymbol{s})^1[i], \beta(\boldsymbol{s})^2[i], .., \beta(\boldsymbol{s})^m[i] \tag{1}$$

is well defined and monotonically decreasing to $\beta(\boldsymbol{s})^m[i] = 0$ for some $m \geq 1$ and this sequence identifies every border of $\boldsymbol{s}[1\,..\,i]$. Here, $\beta(\boldsymbol{s})^k[i]$ is the length of the k-th longest border of $\boldsymbol{s}[1\,..\,i]$, for $1 \leq k \leq m$. Sequence (1) identifies the above-mentioned chain of borders. We will also be using the usual notion of the length of the chain of borders. Clearly, the length of the chain in Sequence (1) is m. Also we use the following notion and notations. In Sequence (1), we call $\beta(\boldsymbol{s})^m[i] = 0$ the *last* value and $\beta(\boldsymbol{s})^{m-1}[i]$ the *penultimate* value in the chain. We now present the following interesting facts that will be useful in our algorithm.

Fact 2 *The length of the chain of borders for a Lyndon Border is at most 2.*

Proof. The result follows, because, if a border is a Lyndon Word, then it cannot itself have a border (Observation 1(1)).

□

Fact 3 *Suppose* $\mathcal{L}\beta(\boldsymbol{s})[i] = x$. *Then* $\mathcal{L}\beta(\boldsymbol{s})[x] = 0$.

Proof. Immediate from Fact 2.

□

Now we are ready to propose a straightforward naive algorithm to compute the Lyndon Border Array $\mathcal{L}\beta(\boldsymbol{s})$ for $\boldsymbol{s}[1\,..\,n]$ as follows.
Algorithm *Naive Lyndon Border Array Construction*

```
1: Compute β(s)[1..n]
2: for i = n → 1 do
3:     Find the penultimate value k of the chain of borders for β(s)[i]
4:     if s[1..k] is a Lyndon word then
5:         Set Lβ(s)[i] = k
6:     else
7:         Set Lβ(s)[i] = 0
8:     end if
9: end for
```

The correctness of the algorithm follows directly from Facts 2 and 3. Now we discuss an efficient implementation of the algorithm. When we traverse through the chain of borders, we reach the penultimate value k and then the last value 0. Clearly, all we need is to check for each such chain, whether $\boldsymbol{s}[1\,..\,k]$ is a Lyndon word. So, we are always interested in finding whether $\boldsymbol{s}[1\,..\,k]$ is a Lyndon word where $\beta(\boldsymbol{s})[k] = 0$. At this point the computation of SNLBfP (Section 3) will be applied. To give an example, by Corollary 8, we know that any border-free prefix to the left of SNLBfP is a Lyndon word and any border-free prefix to the right of SNLBfP is non-Lyndon. This gives us an efficient weapon to check the *If* statement of Line 4 of the above algorithm. Finally, as can be seen below, we can make use of a stack data structure along with some auxiliary arrays to efficiently implement the above algorithm. In particular, we simply keep an array $Done[1\,..\,n]$ initially all false, using a stack and the SNLBfP index (say r). The algorithm is presented below:
Algorithm *Efficient Lyndon Border Array Construction.*

```
for i = 1 → n do
    Set Done[i] =FALSE
end for
Compute β(s)[1..n]
Compute SNLBfP using Algorithm SNLBfP. Say, s[1..r] is the SNLBfP.
for i = n → 1 do
    if Done[i] = TRUE then
        continue                                  ▷ i.e., skip what is done below
    end if
    Compute the chain of β(s)[i] and push each onto a stack S.
    Compute the penultimate value k of the chain of β(s)[i].
    while S is nonempty do
        Pop the value j from the stack
        if  k < r  then                           ▷ i.e., s[1..k] is a Lyndon word
            β(s)[j] = k
        else
            β(s)[j] = 0
        end if
        Done[j] =TRUE
    end while
end for
```

Clearly, the time complexity of the above algorithm depends on how many times a chain of borders is traversed. If we can ensure that a chain of borders is never traversed more than once, then the algorithm will surely be linear. To achieve that we use another array $\mathcal{P}val[1\,..\,n]$, initially all set to -1. This is required to efficiently compute the penultimate value of a chain of borders. The difficulty here arises because we may need to traverse a part of a chain of borders more than once through different indices because two different indices of the border array, $\beta(\boldsymbol{s})$, may have the same value. This may incur more cost and make the algorithm super-linear. To avoid traversing any part of a chain of borders more than once we use the array $\mathcal{P}val[1\,..\,n]$ as follows. Clearly, we only need to traverse the chain of borders to compute the penultimate value. So, as soon as we have computed the penultimate value k, for a chain, we store the value in the corresponding indices of $\mathcal{P}val$. To give an example, suppose we are considering the chain of borders $\beta(\boldsymbol{s})^1[i], \beta(\boldsymbol{s})^2[i], .., \beta(\boldsymbol{s})^{m-1}[i], \beta(\boldsymbol{s})^m[i]$, where $\beta(\boldsymbol{s})^m[i] = 0$, and suppose that the penultimate value is k, i.e., $\beta(\boldsymbol{s})^{m-1}[i] = k$. Now suppose further that the indices involved in the above chain of borders are $i = i_1, i_2, .., i_{m-1}, i_m$. Then as soon as we have got the penultimate value, we update $\mathcal{P}val[i_1] = \mathcal{P}val[i_2] = .. = \mathcal{P}val[i_{m-1}] = k$. How does this help? If the same chain or part thereof is reached for computing the penultimate value, we can easily return the value from the corresponding $\mathcal{P}val[1\,..\,n]$ entry, and thus we never need to traverse a chain or part thereof more than once. This ensures the linear running time of the algorithm.

4.1 For Binary Alphabet

We consider here the case of a binary alphabet Σ, where $\Sigma = \{a, b\}$ and $a < b$. Then a binary string $\boldsymbol{s}$ of length $|\boldsymbol{s}| = n$, where $\boldsymbol{s}$ is a non-letter string (i.e., $n > 1$), can be expressed as:

$$\boldsymbol{s}[1\,..\,n] = (a)^{p'}(b)^{q'}(a)^{p_1}(b)^{q_1}\cdots(a)^{p_m}(b)^{q_m} \tag{2}$$

for $0 \leq p', p_1; \ldots; p_m, q', q_1; \ldots; q_m$. Considering our interest in non-empty binary Lyndon strings, which are border-free, we require the stronger condition that $0 < p', q' < n$ and $0 \leq p_j, q_j$. Let $\mathcal{L}$ be the set of binary Lyndon words.

Clearly, Equation. 2 represents the *Run Length Encoding* $RLE(\boldsymbol{s})$ of a binary string $\boldsymbol{s}$. For example, if $\boldsymbol{s} = (a)(b)(aa)(b)(aaa)(bb)(aa)(b)(aa)(b)$, then

$$RLE(\boldsymbol{s}) = (a)^1(b)^1(a)^2(b)^1(a)^3(b)^2(a)^2(b)^1(a)^2(b)^1$$

We are interested in the values of the parameters (exponents) in Equation. 2 $(p', p_1; \ldots; p_m, q', q_1; \ldots; q_m\ )$ which we will maintain in an auxiliary linked list $\mathcal{R}(\boldsymbol{s})$. Note that $\mathcal{R}[0] = p'$ and $\mathcal{R}[1] = q'$. So for the example string $\boldsymbol{s}$ above, $\mathcal{R}(\boldsymbol{s}) = \{1, 1, 2, 1, 3, 2, 2, 1, 2, 1\}$.

Proposition 17. *For a given string $\boldsymbol{s} \in \Sigma^+$, the* Run Length Encoding *$RLE(\boldsymbol{s})$ and subsequently, the list $\mathcal{R}(\boldsymbol{s})$ can be computed in time and space linear in the size of the given string $\boldsymbol{s}$.*

Now we have the following Proposition to check if a given binary word is a Lyndon word or not (using the *Run Length Encoding* of the binary word).

Proposition 18. *Let $\mathcal{L}$ be the set of Lyndon words over an alphabet Σ, where $\Sigma = \{a, b\}$ and $a < b$. For a non-letter word $\boldsymbol{\ell} = \boldsymbol{\ell}_1\boldsymbol{\ell}_2\cdots\boldsymbol{\ell}_n$ and its corresponding exponents list $\mathcal{R}(\mathcal{L})$, we have the following:*

Case $m = 0$**,** *we have $p_j, q_j = 0$ for $j \in \{1\,..\,m\}$ then $\boldsymbol{\ell} \in \mathcal{L}$ if and only if $p', q' > 0$.*

Case $m > 0$**,** *we have $p', q', p_j, q_j > 0$ then*

1. *$\boldsymbol{\ell} \in \mathcal{L}$ if and only if $p' > p_j$ for $j \in \{1\,..\,m\}$.*
2. *if $p' = p_m$ and $p' > p_j$ for $j \in \{1\,..\,m-1\}$ then $\boldsymbol{\ell} \in \mathcal{L}$ if and only if $q' < q_m \wedge \lambda < 0$.*
3. *if there exists p_j such that $p' = p_j$ and $p' > p_m$ then $\boldsymbol{\ell} \in \mathcal{L}$ if and only if $q' \leq q_j \wedge \lambda < 0$ for $j \in \{1\,..\,m-1\}$.*

where λ is the index of Lyndon invalid point i.e. $\mathcal{L}$-fail condition (As defined in Lemma. 11 and Lemma. 12).

Lemma 19. *For a given string $\boldsymbol{s}$ of length $n = |\boldsymbol{s}|$, let the list $\mathcal{R}$ be the list of exponents of* Run Length Encoding *of $\boldsymbol{s}$, using the list $\mathcal{R}$ it can be checked for each proper prefix of $\boldsymbol{s}$ if it is a Lyndon word or not (computing the array $\Psi(\boldsymbol{s})$) in $O(n)$ time .*

Proof. (**Lemma 19**) We first focus on computing the i-th element $\Psi[i](\boldsymbol{s})$ of the array $\Psi(\boldsymbol{s})$.

Clearly, by Proposition. 18, it takes constant number of comparisons to determine whether the prefix $\boldsymbol{s}[1\,..\,i]$ is a Lyndon word or not.

Since we have to repeat the steps in Proposition.18 for n prefixes, therefore $\Psi(\boldsymbol{s})$ can be done in linear time once $\mathcal{R}(\boldsymbol{s})$ is computed.

□

While computing the border array $\beta(\boldsymbol{s})$ (using the Algorithm in [Lot05]), we can check conditions (1), (2) & (3) in Observation. 1 in constant time, by checking if $\boldsymbol{\ell}_1 < \boldsymbol{\ell}_n$ and $\beta[|\boldsymbol{\ell}|] \leq 0$.

We can check condition (4) (Observation. 1) by computing the parameters according to Proposition. 18 and checking the $\mathcal{L}$-fail conditions (Lemma.11 & 12).

So, we will evaluate the parameters $(p', p_1; \ldots; p_m, q', q_1; \ldots; q_m)$ while computing the lists $\beta[i]$, $\mathcal{R}$ and $\mathcal{L}\beta[i]$ simultaneously.

Note that the i-th entry of *Lyndon Border Array* $\mathcal{L}\beta[i]$ correspond to one of the previous borders (that is also a Lyndon word) and hence we need to consult the value $\Psi[\beta[i]]$, we will set $\mathcal{L}\beta[i] = \beta[i]$ if $\Psi[\beta[i]] = true$. Otherwise, the algorithm recursively finds the longest border in the chain of borders (at position i) that is also a Lyndon word.

An outline of the algorithm is as follows:

1. Compute the border array $\beta(\boldsymbol{s})$.
2. Compute the *Run Length Encoding* and subsequently the list $\mathcal{R}(\boldsymbol{s})$.
3. Compute $\Psi(\boldsymbol{s})$ (Proposition. 18).
4. Check for each prefix of the input string whether or not there exists a border of length $b \neq 0$ such that $\boldsymbol{s}[1\,..\,b] \in \mathcal{L}$.

A binary Lyndon word can also be expressed in terms of Lyndon properties of the integer parameters (exponents) given by its *Run Length Encoding*. For a binary string $\boldsymbol{\ell}$, let $\mathcal{R}(\boldsymbol{\ell_a})$ denote the encoding (as a string) of the subsequence of $\boldsymbol{\ell}$ consisting of all letters a but no letter b $(p', p_1; \ldots; p_m)$, similarly $\mathcal{R}(\boldsymbol{\ell_b})$ denotes the encoding of the subsequence of $\boldsymbol{\ell}$ consisting of all letters b $(q\prime, q_1; \ldots; q_m)$.

Lemma 20. *Let $\boldsymbol{\ell}$ be a binary word with $\boldsymbol{\ell}[1] = a$, $\boldsymbol{\ell}[n] = b$ and associated encodings $\mathcal{R}(\boldsymbol{\ell_a})$ and $\mathcal{R}(\boldsymbol{\ell_b})$. Then $\boldsymbol{\ell}$ is a Lyndon word if and only if, either*
(i) $\mathcal{R}(\boldsymbol{\ell_a})$ is a Lyndon word on the alphabet $\{1 > 2 > 3 > \cdots\}$, or
(ii) $\mathcal{R}(\boldsymbol{\ell_a})$ is a repetition of a Lyndon word as in (i) and $\mathcal{R}(\boldsymbol{\ell_b})$ is a Lyndon word on the alphabet $\{1 < 2 < 3 < \cdots\}$.

Proof. (**Lemma 20**) First we consider necessity. So suppose that (i) holds. Consider any rotation $\boldsymbol{\ell_a^r}$ of $\boldsymbol{\ell_a}$ (including those with split runs of a's). Then $\mathcal{R}(\boldsymbol{\ell_a}) < \mathcal{R}(\boldsymbol{\ell_a^r})$ in lexorder over $\{1 > 2 > 3 > \cdots\}$. Now suppose that (ii) holds. Then for a rotation $\boldsymbol{\ell_a^r}$ of $\boldsymbol{\ell_a}$, either $\mathcal{R}(\boldsymbol{\ell_a}) < \mathcal{R}(\boldsymbol{\ell_a^r})$ in lexorder over $\{1 > 2 > 3 > \cdots\}$, or $\mathcal{R}(\boldsymbol{\ell_a}) = \mathcal{R}(\boldsymbol{\ell_a^r})$ and $\mathcal{R}(\boldsymbol{\ell_b}) < \mathcal{R}(\boldsymbol{\ell_b^r})$ in lexorder over $\{1 < 2 < 3 < \cdots\}$.

For sufficiency, the conditions guarantee that $\boldsymbol{\ell} \in \mathcal{L}$.

□

Run time analysis (For Binary Alphabet) Computing $\beta(\boldsymbol{s})$ can be done in linear time Proposition. 2. Also, *run length encoding* (hence the list $\mathcal{R}$ (the exponents $p', p_1; \ldots; p_m, q', q_1; \ldots; q_m$)) Proposition. 17 and the array $\Psi(\boldsymbol{s})$ can be done in linear time according to Proposition. 18 and Lemma. 19 respectively.

Consequently *Lyndon Border Array* can be computed in linear time once $\beta(\boldsymbol{s})$, $\mathcal{R}(\boldsymbol{s})$ $\Psi(\boldsymbol{s})$ is computed.

The algorithm requires extra space for keeping the arrays $\Psi(\boldsymbol{s})$ and $\beta(\boldsymbol{s})$ (each of length n). Additionally, the algorithm uses an auxiliary extra space to maintain the list $\mathcal{R}(\boldsymbol{s})$ of length at most n. Therefore the total space required to run the procedure is linear to the length of the given string.

Similar time/space efficiency can be achieved to compute *Lyndon Border Array* by simply applying Duval's algorithm [Duv83] on the exponents of a's and then if necessary on those for b's (the concept presented in Lemma. 20).

Lemma 20 shows that a binary string can be decomposed into its Lyndon factors with a double application of Duval's linear algorithm [Duv83], first on the exponents of a and then if necessary on those for b.

5 Lyndon Suffix Array Computation

The well-known suffix array of a string records the lexicographically sorted list of all of its suffixes. Our next contribution is to show how the *Lyndon Suffix Array*, like the original suffix array, can be constructed in linear time; for a string of length n it follows that the indexes in the Lyndon variant will be a subset of $\{1, 2, \ldots, n\}$. We will exploit the elegant fact that Lyndon suffixes are nested:

Fact 4 *If the given string $\boldsymbol{s}$ is a Lyndon word, by Lyndon properties of Lyndon suffixes, the indexes in the Lyndon Suffix Array will necessarily be increasing.*

In order to efficiently construct the Lyndon suffix array we could directly modify the linear-time and space efficient method of Ko and Aluru [KA03] given for the original data structure – this would involve lex-extension ordering (lex-order for substrings, see [DS14]) along with Fact 4. We note that the Ko-Aluru method has also recently been adapted to non-lexicographic V-order and V-letters, and applied in a novel Burrows-Wheeler transform [DS14], and hence is quite a versatile technique. – since the modification here will be similar we will just outline the main steps.

Let an *L-letter* $\boldsymbol{\ell} = \ell_1\ell_2 \ldots \ell_m$ substring denote the simple case of a Lyndon word such that $\ell_1 < \ell_i$ for $2 \leq i \leq m$, assumed to be of maximal length, that is, $\ell_1 = \ell_{m+1}$ (if ℓ_{m+1} exists); hence $|\boldsymbol{\ell}| \geq 1$.

Since, apart from the last letter, there may not be Lyndon suffixes of a string, we perform a linear scan to record the locations of the minimal letter ℓ_1, say, in the string $\boldsymbol{s}$. Observe also that either an L-letter is a Lyndon suffix, or it is the prefix of a Lyndon suffix - the point is that an L-letter is a well-defined chunk of text, a substring of the input $\boldsymbol{s}$, as opposed to the classic single letter approach. In order to sort chunks of text lexicographically, we will apply *lex-extension* order defined as follows.

Definition 21. *Suppose that according to some factorization $\mathcal{F}$, two strings $\boldsymbol{u}, \boldsymbol{v} \in \Sigma^+$ are expressed in terms of nonempty factors: $\boldsymbol{u} = u_1 u_2 \cdots u_m, \boldsymbol{v} = v_1 v_2 \cdots v_n$. Then $\boldsymbol{u} <_{LEX(F)} \boldsymbol{v}$ if and only if one of the following holds:*
(1) $\boldsymbol{u}$ is a proper prefix of $\boldsymbol{v}$ (that is, $u_i = v_i$ for $1 \leq i \leq m < n$); or
(2) for some $i \in 1..min(m,n), u_j = v_j$ for $j = 1, 2, .., i-1$, and $u_i < v_i$ (in lexicographic order).

First, using a linear scan we apply the $\mathcal{L}\beta$ to record the indexes of all the Lyndon suffixes of $\boldsymbol{s}$. The factorization $\mathcal{F}$ which we will use is that of decomposing the input into substrings of L-letters; in the case of unit length L-letters we concatenate them so that $\mathcal{F}$ has the general form:

$$\boldsymbol{u}\ell_1^{i_1}\boldsymbol{\ell_1}\ell_1^{i_2}\boldsymbol{\ell_2}\cdots\ell_1^{i_t}\boldsymbol{\ell_t}$$

where $\boldsymbol{u} \in \Sigma^*$ does not contain ℓ_1, $t \geq 1$, each $i_j \geq 0$, every $\boldsymbol{\ell_j}$ is an L-letter and $|\boldsymbol{\ell_t}| \geq 1$. – in practice this just entails keeping track of occurrences of the letters $\boldsymbol{\ell}_1$ in $\boldsymbol{s}$. Since we are computing Lyndon suffixes we can ignore $\boldsymbol{u}$ and hence assume that $\mathcal{F}$ has the form $\ell_1^{i_1}\boldsymbol{\ell_1}\ell_1^{i_2}\boldsymbol{\ell_2}\cdots\ell_1^{i_t}\boldsymbol{\ell_t}$.

An ℓ-suffix is a suffix of $\mathcal{F}$ commencing with $\ell_1^{i_j}$. We now apply the Ko-Aluru linear method, consisting of three main steps, to the ℓ-suffixes: which we first outline followed by further detail for each.

- Using a linear scan of the input string $\boldsymbol{s}$ and lex-extension ordering, divide all ℓ-suffixes of $\boldsymbol{s}$ into two types: That is, let $\boldsymbol{s_i}$ denote the suffix starting at index i, so $\boldsymbol{s_i} = \boldsymbol{s}[i..n]$. Then the type S Lyndon suffixes are the set $\{\boldsymbol{s_i} | \boldsymbol{s_i} < \boldsymbol{s_{i+1}}\}$ and the type L Lyndon sufixes are the set $\{\boldsymbol{s_j} | \boldsymbol{s_j} > \boldsymbol{s_{j+1}}\}$.
- Sort by lex-extension all ℓ-suffixes of type S in $O(n)$-time using a modified Bucket Sort followed by recursion on at most half of the string.
- Using a linear scan obtain the lex-extension order of all remaining ℓ-suffixes (assumed to be type L) from the sorted ones. This step is obtained from observing that the type L ℓ-suffixes occurring in $\boldsymbol{s}$ between two type S ℓ-suffixes, S_i and S_j where $i < j$, are already ordered such that $L_l <_{LEX(\mathcal{F})} L_k$ for $i < k < l < j$.

At this stage we have computed an ℓ-suffix array. There are two final steps for computing the Lyndon suffix array. Firstly, the classic suffix array for the last L-letter $\boldsymbol{\ell_t}$ (without the prefix ℓ_1) is processed directly using the Ko-Aluru method and the indexes inserted into the ℓ-suffix array. Then we perform a linear scan of the ℓ-suffix array, and applying Fact 4, by selecting only the sequence of increasing integers yields the Lyndon suffix array.

The last suffix is both type S and L. The modification from lexicographic to lex-extension order follows from, firstly the linear factorization of the input into L-letters, and secondly that lex-extension ordering applies lexicographic order pairwise to L-letter substrings which each requires no more than time linear in the length of the L-letters – hence $O(n)$ overall. The space efficiency follows from the original method [KA03].

5.1 A Simpler algorithm for computing a Lyndon Suffix Array from a Suffix Array

We present an alternative simple algorithm, derived from the classic suffix array, which also exploits the nested structure expressed in Fact 4. Suppose we are given the suffix array of the string $\boldsymbol{s}$. Now our algorithm finds the largest suffix (max), and then searches inside it to find the second largest one and so on, taking advantage of the fact that the suffixes are already sorted in the suffix array. Repeatedly finding the max value can be implemented efficiently using the Range Minimum Query (RMQ) [BFC00], which requires $O(n)$ time preprocessing and then $O(1)$ time for each query.

Lemma 22. ***Maximum range suffixes are Lyndon words:*** *Suppose we are given the suffix array $\mathcal{A}$ of a string $\boldsymbol{s}$. Let the set $\mathcal{M}$ be the set of maximum values of the range of suffixes $\mathcal{A}[0\,..\,i]$, where i is the index of the i-th suffix ℓ_i. Each suffix $\ell \in \mathcal{M}$ is a Lyndon word.*

Proof (Lemma 22). Suppose ℓ is a max range suffix ($\ell \in \mathcal{M}$) at index i with order value $SA[i] = m$. Suppose there exists a suffix $\boldsymbol{\ell}''$ of $\boldsymbol{\ell}$ at index i'' (w.r.t. to $\mathcal{A}$) with order value $\mathcal{A}[i''] = m''$, and also suppose $\boldsymbol{\ell}'' < \boldsymbol{\ell}$. Hence $m'' > m$ ($\mathcal{A}[i''] > \mathcal{A}[i]$) where $i'' < i$, which contradicts the fact that m is the maximum value in the range $\mathcal{A}[0\,..\,i]$. Therefore, $\boldsymbol{\ell}$ is strictly smaller than all of its proper suffixes – we conclude that ℓ is a Lyndon word.

□

The linear-time method for computing the Lyndon suffix array is very simple. Below, we first outline the steps followed by the pseudo-code.

1. Compute the suffix array $\mathcal{A}$ of $\boldsymbol{s}\$$.
2. Find the value $max = \mathrm{MAX}(\mathcal{A}[0\,..\,n])$ and its index i in the suffix array, then add the value max to Lyndon Suffix array $\mathcal{L}S$.
3. Find the value max and its index i in the range $\mathcal{A}[0\,..\,i]$, $max = \mathrm{MAX}(\mathcal{A}[0\,..\,i])$ and $i = \mathrm{Indexof}(max, \mathcal{A})$, then add the value max to $\mathcal{L}S$.
4. Repeat step 2 until the set $\mathcal{A}[0\,..\,i]$ is empty.

Algorithm 1 Computing the Lyndon Suffix Array from a Suffix Array.

```
procedure ComputeLSA(s)
    n ← |s|;  L[1..n] ← (−1)^n
    ▷ compute the suffix array of s$
    A ← SA(s$)
    i ← n; j ← n; max ← 0
    while (A[0..i] ≠ ∅) do
        ▷ find the maximum value (and its index) in the range A[0..i].
        (max, i) ← (A[idx], idx) | idx = arg max_idx(A[idx]) for 0 ≤ idx < i
        j ← j − 1
        L[j] ← max
    end while
    return L
end procedure
```

6 Conclusion

In this article, we have extended two well-known data structures in stringology. We first have adapted the concept of a border array to introduce the ***Lyndon Border Array*** $\mathcal{L}\beta$ of a string $\boldsymbol{s}$, and have described a linear-time and linear-space algorithm for computing $\mathcal{L}\beta(\boldsymbol{s})$. Furthermore, we have defined the ***Lyndon Suffix Array***, which is an adaptation of the classic suffix array. By modifying the linear-time construction of Ko and Aluru [KA03] we similarly achieve a linear construction for our Lyndon variant. We also present a simpler algorithm to construct a ***Lyndon Suffix Array*** from a given Suffix Array.

The potential value of the Lyndon Border Array is that it allows for deeper burrowing into a string to yield paired Lyndon patterned substrings. The Lyndon suffix array lends itself naturally to searching for Lyndon patterns in a string. If the given text or string has a sparse number of Lyndon words (as likely in English literature due to the vowels a, e often occurring internally in words), then the Lyndon suffix array may offer efficiencies. Polyrhythms, or cross-rhythms, are when two or more independent rhythms play at the same time – nested Lyndon suffixes can exist in these rhythms. We propose that applications of these specialized data structures might arise in the context of the relationship existing between de Bruijn sequences and Lyndon words [FM78].

We propose that applications of ***Lyndon Border Array*** may arise in combinatorics in relation to the Christoffel words. The methods used for these problems often make use of structures equivalent to suffix trees in order to achieve efficient execution.

References

[AC95] Alberto Apostolico and Maxime Crochemore. Fast parallel lyndon factorization with applications. *Mathematical Systems Theory*, 28(2):89–108, 1995.

[BFC00] Michael A. Bender and Martin Farach-Colton. The lca problem revisited. In Gaston H. Gonnet, Daniel Panario, and Alfredo Viola, editors, *LATIN*, volume 1776 of *Lecture Notes in Computer Science*, pages 88–94. Springer, 2000.

[CFL58] K. T. Chen, R. H. Fox, and R. C. Lyndon. Free differential calculus, iv — the quotient groups of the lower central series. pages 68:81–95, 1958.

[Che04] M. Chemillier. Periodic musical sequences and Lyndon words. *Soft Comput.*, 8(9):611–616, 2004.

[CHL07] Maxime Crochemore, Christophe Hancart, and Thierry Lecroq. *Algorithms on strings*. Cambridge University Press, 2007.

[CR02] Maxime Crochemore and Wojciech Rytter. *Jewels of stringology*. World Scientific, 2002.

[DR04] O. Delgrange and E. Rivals. STAR: an algorithm to search for tandem approximate repeats. *Bioinformatics*, 20(16):2812–2820, 2004.

[DS14] J. W. Daykin and W. F. Smyth. A bijective variant of the Burrows-Wheeler transform using V-order. *Theoret. Comput. Sci.*, 531:77–89, 2014.

[Duv83] J. P. Duval. Factorizing words over an ordered alphabet. *J. Algorithms*, 4(4):363–381, 1983.

[FM78] H. Fredricksen and J. Maiorana. Necklaces of beads in k colors and k-ary de Bruijn sequences. *Discrete Mathematics*, 23(3):207–210, 1978.

[GS12] J. Y. Gil and D. A. Scott. A bijective string sorting transform. *CoRR*, abs/1201.3077, 2012.

[KA03] P. Ko and S. Aluru. Space efficient linear time construction of suffix arrays. In Ricardo A. Baeza-Yates, Edgar Chávez, and Maxime Crochemore, editors, *CPM*, volume 2676 of *Lecture Notes in Computer Science*, pages 200–210. Springer, 2003.

[Lot83] M. Lothaire. *Combinatorics on Words*. Reading, MA (1983); 2nd Edition, Cambridge University Press, Cambridge (1997). Addison-Wesley, 1983.

[Lot05] M. Lothaire. *Applied Combinatorics on Words (Encyclopedia of Mathematics and its Applications)*. Cambridge University Press, New York, NY, USA, 2005.

[Lyn54] R. C. Lyndon. On burnside problem *i*. *Transactions of the American Mathematical Society*, 77:202 – 215, 1954.

[Lyn55] R. C. Lyndon. On burnside problem *ii*. *Transactions of the American Mathematical Society*, 78:329 – 332, 1955.

[MP70] J.H. Morris and Vaughan R. Pratt. A linear pattern-matching algorithm. *Technical Report 40*, (8), 1970.

[NS13] Shoshana Neuburger and Dina Sokol. Succinct 2d dictionary matching. *Algorithmica*, 65(3):662–684, 2013.

[Per05] L. Perret. A chosen ciphertext attack on a public key cryptosystem based on Lyndon words. *IACR Cryptology ePrint Archive*, 2005:14, 2005.

[Reu93] C. Reutenauer. *Free Lie algebras*. London Math. Soc. Monographs New Ser. 7. Oxford University Press, 1993.

[Smy03] B. Smyth. *Computing Patterns in Strings*. ACM Press Bks. Pearson/Addison-Wesley, 2003.